REMARKS/ARGUMENTS

This Amendment is responsive to the Office Action mailed on December 6, 2007. In this Amendment, claims 1-3, 5-8, and 10-24 are pending and subject to examination.

I. Objections to the Claims

Claims 1 and 10 are objected to because of certain informalities contained in the claims that have been interpreted as claiming intended use. Claims 1 and 10 have been amended to cure the informalities. Applicants request that the objections be withdrawn.

II. Rejections under 35 U.S.C. §112

Claims 1-24 are rejected under 35 U.S.C. §112 as failing to comply with the written description requirement. Applicants respectfully submit that there is ample support in the Specification for all of the claims.

In particular, the rejection is based on the assertion that the following claim language in claims 1 and 10 lacks support in the Specification: "such that at least some of the replication storage volumes are located outside the respective failure boundary ***for any of the types of storage failure.***" [Emphasis added by the Examiner]

The Specification gives many examples of the types of the storage failure that are at issue in the present application. For example, paragraphs 29-30 in the Specification discuss failure boundaries such as ECC groups, controller pair groups, subsystem groups, logical volume groups and volume groups. One skilled in the art would recognize that there are many possible failure boundary schemas that could be implemented in a given embodiment. The key disclosure is that failure boundaries are used to designate all of the volumes affected by the failure used to define the boundary (Par. 29).

The Specification also describes how secondary volume groups are distributed in a way that crosses these defined failure boundaries. For example, paragraph 31 discloses that "[a]s mentioned, preferably the [secondary] groups will cross failure boundaries. The level of the failure boundary will be determined automatically using the system software and an

appropriate policy, or the level may be determined by a system administrator. However determined, the level can consist of one error correction group 301, a controller pair 302, etc., as discussed above." The secondary groups discussed in this passage are one possible embodiment for the replication storage volume claimed. An illustration of how data can be stored across these secondary groups is illustrated in Figure 4.

These two passages make it clear that if any storage failure used to define a failure boundary occurs, then at least some of the secondary groups will be located outside of the failure boundary. This is true because the secondary groups have been distributed in a way that crosses these failure boundaries.

The Applicants have shown that there is plenty of support in Specification for the claim language. Claims 1 and 10 have also been amended to clarify the emphasized claim language "for any types of storage failure." The storage failures referenced in this passage are the same failures used to determine the plurality of failure boundaries. Applicants respectfully submit that 35 U.S.C. §112 rejection for claims 1-24 be withdrawn.

III. Rejections under 35 U.S.C. §103

Claims 1-20 and 23-24 are rejected under 35 U.S.C. §103(a) as being obvious over Bridge (US 6,530,035) in view of Iwanmi (2002/0112030) and Ohran (US 2002/0112134). Applicants respectfully submit that these references do not teach or suggest each element of these claims.

In particular, Ohran does not teach "a first type of content to be stored has replication storage volumes assigned across each failure boundary, such that at least some of the replication storage volumes are located outside the respective failure boundary for any of the types of storage failure." Ohran teaches that *different* types of data can be stored in *different* failure boundaries, but it does not teach that a single type of data can be stored across storage boundaries.

As the Examiner has pointed out, Ohran discusses two different types of data that are stored to facilitate data backups: incremental backup data stored in preservation memory and full backup data stored in a mass storage device.

Ohran teaches that the mass storage device generally is "a writable, nonvolatile mass storage device. In particular, mass storage device 12 can be the hard drive associated with a conventional personal computer or any other storage volume that is used to store data obtained from computer 10." (Page 3, Par. 28) Ohran also teaches that preservation memory "can be a volatile device, such as random access memory (RAM) or any other device that can store data blocks that are to be overwritten in mass storage device 12. Although preservation memory 14 is illustrated as being a separate device in Fig. 1, the preservation memory can be a partition or another portion of mass storage device 12." (Page 3, Par. 29)

In these sample embodiments taught by Ohran, if the mass storage device fails, then all of the full backup data would be lost. It is unknown whether the incremental backup data would also be lost as a result of the same failure without more information. If the preservation memory storing the incremental backup data was in fact in a different failure boundary than the mass storage device failure, then the incremental backup data would still be available. Otherwise, if the preservation memory did reside in the same failure boundary as the mass storage device, then the incremental data would be fully lost as well. Ohran is effectively teaching that these different types of data, incremental backup data and full backup data, can be independently stored in different failure boundaries.

The teaching in Ohran contrasts with what is claimed in the present invention. Ohran does not teach how to store a given type of content across failure boundaries. Storing data across a failure boundary affects what portions of the content is lost when a failure event occurs.

As in the Office Action, the full backup data and incremental backup data in Ohran can be treated as the first and second types of content in the Applicants' claims, respectively. As set forth in the Applicants' claims, a first data type can be stored in replication storage volumes so that "at least some of the replication storage volumes are located outside the respective failure boundary for any of the types of storage failure." If the full backup data taught in Ohran is treated as the first content type in the Applicants' claims and is stored as claimed, then only a subset of the full backup data would be lost for any failure condition used to define a failure boundary. In Ohran, all of the data could be lost due to a failure condition such as the mass storage device failing. If the failure of the mass storage device is used as a failure

boundary as in the Applicants' claims, then only a subset of the full backup data would be lost in due to that failure event. Ohran teaches steps that can be taken to try to recover from this failure, but the steps taken after a failure event are not the subject matter of the Applicants' claims. The Applicants' claims deal with how data is stored before a failure event occurs.

Similarly, if the incremental data is treated as the first content type in the Applicants' claims and is stored as claimed and the full backup data is treated as the second content type, then only a subset of the incremental data would be lost for any failure condition used to define a failure boundary. In Ohran, the data that is lost as a result of a failure event is necessarily an all or nothing situation with respect to each data type because Ohran does not teach how to store a single data type across a failure boundary.

Claim 23 & 24 point out the distinction with Ohran even more clearly because they define similar types of data to those used in Ohran. In these dependent claims, the first data type is defined as full backup data and it is stored in a way that crosses failure boundaries. The second data type is defined as differential backup data that may be stored within a single failure boundary.

In claim 23 or 24, if a full backup of five pieces of data is made (e.g., A, B, C, D, and E as in Ohran), then at least some of this full backup data would survive a failure used to define a failure boundary, because the data would be assigned to replication volumes that went across this failure boundary. Staying with this scenario, it is possible, although not required, that all of the incremental backup data (e.g., A, D, D_1, B, and D_x,as in Ohran) can reside within a single failure boundary. As an example reflecting Applicants' claims, say that a storage system stores A, B, and C from the full backup data in one failure boundary, and D and E from the full backup data and all of the incremental backup data in another failure boundary. Of course, the data could be divided up further into more than two groups, but two groups will suffice for this example. If the failure condition used to define this failure boundary occurs, then either A, B, and C or D, E and the incremental backup data would survive this failure event. Some of the full backup data will always survive this failure event and potentially all of the preservation memory could survive as well. If this data was stored as is taught in Ohran, then either all of the full backup data would be lost, all of the incremental backup data would be lost, or all of the data

would be lost. Since Ohran does not teach how to store different content types across failure boundaries, the possible outcomes of a give storage failure are very different than what is possible with the claimed invention.

One rationale for treating these two data types in this way is that differential or incremental backup data is not as useful in the absence of a full set of original data, and so there is not as great a need to store this data across failure boundaries. As recited in Applicants' claims, a user is able to use the storage area more efficiently and flexibly by allowing the user to dictate how data is stored either across or within failure boundaries based on the type of data at issue. In Ohran, while the full backup data and incremental backup data can reside in different failure boundaries, there is no scenario where the user can define how a given set of this data is stored across failure boundaries based on the data type.

Ohran does teach that one type of data that resides in one failure boundary, such as incremental backup data in preservation memory, can be used to rebuild a different type of data stored in a different failure boundary, such as full backup data in a mass storage device. In essence, Ohran is teaching one way that data can be used across failure boundaries. However, this feature is not what is claimed in the present application. The claims of this application are targeted at how different types of data can be distributed across failure boundaries. How data is used after being stored is not what is currently being claimed.

Bridge and Iwami also do not teach or suggest that "a first type of content to be stored has replication storage volumes assigned across each failure boundary, such that at least some of the replication storage volumes are located outside the respective failure boundary for any of the types of storage failure," nor do they teach "a second type of content is able to be stored having replication storage volume within at least one failure boundary." Because Bridge, Iwami, and Ohran, cannot be combined in a way that teaches or suggests all of the features of the claims 1 and 10, and since all of the other claims are dependent on these two independent claims, Applicants respectfully request that the obviousness rejections with respect to these claims be withdrawn.

Claims 21-22 are also rejected under 35 U.S.C. §103(a) as being obvious over Bridge (US 6,530,035) in view of Iwanmi (2002/0112030) and Ohran (US 2002/0112134). Claims 21-22 are dependent on claims 1 and 10, and so claims 21-22 should be allowed for the same reasons as those independent claims. Additionally, the Office Action cites two passages from the Applicants' Specification in support of the §103(a) rejection. The Office Action states that Bridge and Iwami do not disclose "the primary storage volumes and replication storage volumes are horizontally or vertically addressed." Applicants note that it is improper to then rely on the Applicants' own Specification as a basis for teaching this element. *See* MPEP 2142 (stating that knowledge of the applicant's disclosure must be put aside when making a determination of obviousness). Consequently, Applicants respectfully submit that this rejection to claims 21-22 be withdrawn.

Accordingly, the limitations of claims 1-24 are not taught or suggested by Bridge in view of Iwanmi and Ohran. Applicants thus respectfully request that the rejection to claims 1-24 be withdrawn.

CONCLUSION

In view of the foregoing, Applicants believe all claims now pending in this Application are in condition for allowance. The issuance of a formal Notice of Allowance at an early date is respectfully requested.

If the Examiner believes a telephone conference would expedite prosecution of this application, please telephone the undersigned at 925-472-5000.

Respectfully submitted,

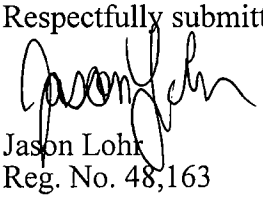

Jason Lohr
Reg. No. 48,163

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, Eighth Floor
San Francisco, California 94111-3834
Tel: 925-472-5000
Fax: 415-576-0300

Attachment
JDL:JRC:slh
61338674 v1